

2025 RELEASE UNDER E.O. 14176

ARBITRATION IN LOCAL SYSTEM FOR ACCESS  
TO MEMORY IN A DISTANT SUBSYSTEM

Patrick J. Smith

TECHNICAL FIELD OF THE INVENTION

The technical field of this invention is data movement in multiprocessor systems.

5 BACKGROUND OF THE INVENTION

Microprocessor systems employing multiple processor subsystems including a combination of local and shared memory are becoming increasingly common. Such systems normally have interconnect formed in large part by wide busses carrying data and control information from one subsystem to another.

Busses are at one instant of time controlled by a specific module that is sending information to other modules. A classical challenge in such designs is providing bus arbitration that guarantees that there are no unresolved 5 collisions between separate modules striving for control of the bus.

#### SUMMARY OF THE INVENTION

The preferred embodiment of this invention relates to bus 10 arbitration in a Multiple-DSP Shared-Memory (MDSM) systems. The preferred embodiment MDSM contains four fixed point DSP cores and a total of 896K Words of on-chip single-access RAM (SARAM) and dual-access RAM (DARAM). It is highly optimized for remote access server (RAS) or remote access concentrator 15 (RAC) and other DSP applications.

This invention comprises an arbitration technique for bus access in a multiple DSP system having four-way shared DARAM memory modules. A DARAM4W Wrapper envelops and includes the shared DRAM memory. It includes all the necessary arbitration 20 and data steering logic to resolve simultaneous access requests by four program "read" ports, the local peripheral port and the local program "write" port.

In each DARAM up to two accesses can occur every clock cycle, one on each one-half clock period. The ports are 25 hardwired to a particular one-half cycle for simplicity of operation. This maintains a one wait state requirement for the design under normal operating conditions. Arbitration among the four local DARAM selects, peripheral bus (M bus) writes and program writes is performed in the DARAM4W Wrapper. 30 A global traffic module decodes, in straightforward fashion,

all input program page addresses and generates the four local DARAM selects. Arbitration between the two simultaneous program page accesses to the neighbor DARAM is performed within the global traffic module.

5

#### BRIEF DESCRIPTION OF THE DRAWINGS

These and other aspects of this invention are illustrated in the drawings, in which:

Figure 1 illustrates in high level block diagram form a  
10 multiple DSP, shared memory (MDSM) system;

Figure 2 illustrates the individual functional blocks of one subsystem of an MDSM system;

Figure 3 illustrates in high level block diagram form the DARAM4W wrapper of representative subsystem A;

15 Figure 4 illustrates the address set-up time for the first half access, a full clock period; and

Figure 5 illustrates the address set-up time for the second half access, only a one-half clock period.

20 DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

The present invention relates to bus arbitration in a Multiple-DSP Shared-Memory (MDSM) system. The MDSM system of the preferred embodiment contains four fixed point DSP cores and a total of 896K Words of single-access RAM (SARAM) and  
25 dual-access RAM (DARAM). A high-level block diagram of this MDSM system is illustrated in Figure 1. The four subsystems A 101, B 102, C 103 and D 104, are each connected to the other subsystems via four read busses entering the bus switching networks 100 at locations 116, 136, 156 and 176.

DSP core 111 of subsystem A 101 accesses shared memory 153 in subsystem C 103 by way of its global traffic module 115. DSP core 111 also accesses shared memory 133 in subsystem B 102 and shared memory 173 in subsystem D, both by way of 5 global traffic module 135 of subsystem B 102. The subsystems C 103 and D 104 are "far" subsystems to subsystem A 101. This means that propagation delays are longer for such accesses than for "close" accesses. Each DSP core such as DSP core 111 includes data manipulation, data access and program flow 10 control hardware. The data manipulation hardware typically includes: an integer arithmetic logic unit (ALU); a multiplier, which may be part of a multiply-accumulate (MAC) unit; a register file including plural data registers; and may include special purpose accelerator hardware configured for 15 particular uses. The data access hardware typically includes: a load unit controlling data transfer from memory to a data register within the register file; and a store unit controlling data transfer from a data register to memory. Control of data transfer by a load unit and a store unit 20 typically employs address registers storing the corresponding memory addresses as well as address manipulation hardware such as for addition of the contents of an address register and an index register or immediate field. DSP core 111 may include plural units of each type and operate according to superscalar 25 or very long instruction word (VLIW) principles known in the art. The program flow control hardware typically includes: a program counter storing the memory address of the current instruction or instructions; conditional, unconditional and calculated branch logic; subroutine control logic; interrupt 30 control logic; and may also include: instruction prefetch

logic; and branch prediction logic. The exact structure of DSP core 111 is not as important as that it functions as a computer central processing unit.

Paths 190 leading from subsystem C 103 shared memory 152 and subsystem D 104 shared memory 173 to DSP core 111 illustrates symbolically such a "far" path. Subsystem B is a "close" subsystem to subsystem A 101. This means that propagation delays are shorter for such accesses than for "far" accesses. Path 195 leading from subsystem B 102 shared memory 133 to DSP core 111 illustrates symbolically such a "close" path.

Each subsystem has a corresponding set of "close" and "far" access paths for its own DSP. The "program read" cycle in which such "read" accesses will be performed are selected 15 for the "close" and "far" accesses. Four "program read" accesses are defined. PROGRAM READ 1 and PROGRAM READ 2 are initiated at the beginning of the first half of clock cycle; PROGRAM READ 3 and PROGRAM READ 4 are initiated at the beginning of the last half of clock cycle. Table 1 lists, for 20 each subsystem and DSP, the local, close and far path accesses illustrated in Figure 1.

Subsystem/DSP	Local	Close Paths/Cycle	Far Paths/Cycle
Subs A/111	112	195 A, B READ 3, 4	190 C, D READ 1, 2
Subs B/131	132	A, B READ 3, 4	C, D READ 1, 2
Subs C/151	152	C, D READ 3, 4	A, B READ 1, 2
Subs D/171	172	C, D READ 3, 4	A, B READ 1, 2

Table 1

The MDSM system paths, by which the four-way shared dual access RAM data flows, are directed by way of the global traffic modules (traffic module 114 in subsystem A 101). Each global traffic module drives a four-way shared DARAM wrapper (DARAM4W 127 in subsystem A 101) that contains the arbitration logic necessary to avoid bus collisions.

Figure 2 illustrates in block diagram from individual functional blocks comprising subsystem A 101. Subsystems B 102, C 103 and D 104 are identical to subsystem A 101. DSP core 111 has "read" access within subsystem A 101 to unshared local RAM 112 via bus program (P) bus 130 and shared RAM 113 also via P bus 130. DSP core 111 has "write" access within subsystem A 101 to unshared local RAM 112 and shared RAM 113 via E bus 122. By way of three additional busses 124, 125, and 126, DSP core 111 also has read access to shared RAM outside subsystem A 101 in the other three subsystems B 102, C 103 and D 104. Summarizing, four of the six paths from

subsystem A 101 shared memory which must be arbitrated by the DARAM4W wrapper 113 are: "read" path 116 from shared memory 113 of subsystem A 101 to a DSP core of another of the three subsystems; "read" path 124 from shared memory 133 of 5 subsystem B to DSP core 111 of subsystem A; "read" path 125 from shared memory 153 of subsystem C to DSP core 111 of subsystem A; and "read" path 126 from shared memory 173 of subsystem D to DSP core 111 of subsystem A.

RAM functions for the entire MDSM system are categorized 10 as local memory, four-way shared memory and described as follows. The local memory preferably includes: 512 KW of zero wait state data SARAM, 128 KW per subsystem such as local DARAM and SARAM 112 illustrated in Figure 2; and 128 KW zero wait state data/program DARAM, 32 KW per subsystem such as 15 local DARAM and SARAM 112 illustrated in Figure 2. The four-way shared memory preferably includes: 256 KW one wait state program DARAM shared by subsystems A 101, B 102, C 103 and D 104, 64 KW per subsystem such as four-way shared DARAM4W 113 illustrated in Figure 2.

20 Referring to Figure 2, the traffic module 114 decodes address 108 of DSP P bus 130 and generates control signals 118 that make the memory bank selection between the local memory blocks of local DARAM and SARAM 112. Traffic module 114 also multiplexes the received acknowledge signals and "read" data 25 from the memory blocks to DSP core 111 via lines 119.

The global traffic module 115 decodes the address 109 of DSP P bus 130. Global traffic module 115 drives memory bank selects 117 to the four-way shared memory wrapper 127 and decodes two program address busses 109 to determine if an 30 access is to the local block of global memory or to global

memory associated with another subsystem. Because Figure 2 is describing a particular subsystem (in this case subsystem A 101), there is an additional task its global traffic module must perform. Global traffic module 115 arbitrates access by 5 signals 128 two of the other subsystems to a third subsystem for four-way shared program "read". Finally, it also communicates a global acknowledge signal 129 as part of its communication with DSP core 111.

Each MDSM subsystem contains a DARAM wrapper. DARAM4W 10 113 includes wrapper 127 illustrated in Figure 2. Each DSP core is capable of accessing a 128K word block of four-way shared memory with one wait state. Wrapper 127 interfaces local, close and far accesses to the shared portion of DARAM4W 113, that is the shared 32K Word block of memory. DARAM4W 15 wrapper 127 supports a total of six interfaces: Program READ bus A 130 for DSP access; Program READ bus B 124 for DSP access; Program READ bus C 125 for DSP access; Program READ bus D 126 for DSP access; M read/write bus 121 for peripheral access; and E data write bus 122 for DSP access. The basic 20 function of wrapper 127 is to arbitrate access to the memory among these six interfaces. This involves arbitration for program "reads" among four cores, local peripheral and local program writes contending for two accesses, one on each one-half clock cycle.

Global traffic module 115 decodes the program page address for access to either its local DARAM or a neighbor DARAM. It generates a total of eight memory bank select signals of which four are local. Arbitration between the four local DARAM selects, M bus 121 writes and program writes is 30 performed in wrapper 127. Global traffic module 115 does a

straight forward decode for both input program page addresses, and generates four local DARAM selects. Arbitration between the two conflicting program page accesses to the neighbor DARAM is also performed within the global traffic module 115.

5 Referring again to Figure 1, one can see that the route delay on the acknowledge "ack" signal to subsystem C 103 or subsystem D 104 for access to memory in subsystem A 101 would be unnecessarily long if it was generated from wrapper 127 in subsystem A 101. Instead, the "ack" signal 169 can be  
10 generated by the global traffic module 155 in subsystem C 103. Global traffic module 155 is physically closer to both subsystems C 103 and D 103 minimizing the route delay on the "ack" signal.

Program access to the "far" neighbor DARAM occurs in the  
15 first half of the cycle as these accesses provide a full cycle of setup time on the address. For subsystem A 101 the "far" neighbor DARAMs are those of subsystem C 103 and subsystem D 104. The requesting core is physically furthest from the target DARAM, so a full cycle address set up is required.

20 The local M bus 121 "read" port also competes for first half access. Local M bus 121 "reads" always have priority and are never stalled. Both page accesses to the neighbor DARAM are arbitrated every time both cores make a request simultaneously assuming there are no local M bus 121 requests.

25 Arbitration of conflicts between the two program page accesses to the neighbor DARAM4W 113 is performed within the global traffic module 115. The priority amongst PAGE 1 and PAGE 2 changes every time PAGE 1 and PAGE 2 both request access to the memory on the same cycle. Initially PAGE 1 will  
30 have priority over PAGE 2. A single register bit controls the

priority. If a request from both PAGE 1 and PAGE 2 occurs simultaneously, priority is given to PAGE 1. The PAGE 1 bus request will complete, and the PAGE 2 bus request will be stalled one clock cycle. The priority register will toggle,  
5 so at the next occurrence of a simultaneous request by PAGE 1 and PAGE 2, PAGE 2 will be given top priority. The priority changes only when there is a collision between PAGE 1 and PAGE 2.

Wrapper 127 arbitrates access by the four program "read" ports, the local peripheral port and the local program "write" port. Up to two accesses to the memory can occur every clock cycle. An access is granted on each one-half clock cycle. The ports are hardwired to a particular one-half cycle in order to simplify operation. Table 2 lists the accesses to be made on each half-clock cycle, and identifies the arbitration priority and requirements. The paths for these program "reads", Program READ 1, Program READ 2, Program READ 3, and Program READ 4 were indicated in Table 1 for the reference numbered paths in Figure 1.

20

First Half Cycle	Second Half Cycle
M Bus READ	M Bus Write
Program READ 1, READ 2 toggle	Program Write
	Program READ 3, READ 4 toggle

Table 2

Within a one-half cycle time interval only one of the  
25 possible requesters is granted access to the memory. The

remaining requesters are stalled for one clock by driving a bus acknowledge signal low.

Program reads 1 and 2 contend for the first half of the cycle, while program READS 3 and 4 contend for the second half of the cycle. The address set-up time for the first half access is a full clock period, while the address set up time for the second half is only a half clock period. Table 3 lists the connection paths of the physical memory to the program busses for each of the four subsystems.

10

Subsystem	Physical Memory	READ 1	READ 2	READ 3	READ 4
Subs A	4MP0/4MP1	Prog C	Prog D	Prog A	Prog B
Subs B	4MP2/4MP3	Prog C	Prog D	Prog A	Prog B
Subs C	4MP4/4MP5	Prog A	Prog B	Prog C	Prog D
Subs D	4MP6/4MP7	Prog A	Prog B	Prog C	Prog D

Table 3

The M bus 121 read is always given top priority in the first half cycle. These signals will be serviced immediately and are never stalled. Program "reads" for bus A and B contend for the first half of the cycle, while program "reads" for bus C and D contend for the second half of the cycle. Bus 1 and bus 2 compete for the memory in the first half cycle. Each DARAM4W is wired such that bus 1 and bus 2 are driven from the other half of the chip. That is, DARAM4W 113 in subsystem A has bus 1 connected to Program C and bus 2 connected to Program D. This is done to provide the most distant cores adequate setup time.

The priority between READ 1 and READ 2 toggles every time READ 1 and READ 2 both request access to the memory on the same cycle. This has been previously described. The priority only changes when there is a collision between READ 1 and READ  
5 2. The arbitration logic for the READ 1 and READ 2 busses is contained in the global traffic module of the other half sub-system. The arbitration for access to the four-way DARAM 113 of subsystem A 101 (4MP0/4MP1) is done in the global traffic module of subsystem C 103. This global traffic module  
10 provides the acknowledges to subsystem C 103 and subsystem D 104 for access to memory in subsystem A 101.

This approach minimizes several important parameters. This approach minimizes the propagation delay of the program page address. This minimizes the propagation delay of the  
15 "ack" signal to the requesting subsystem. It minimizes the number of signals between subsystems for four-way memory.

The multiplexing of the program "read" addresses and data for M bus 121 "reads", READ 1 and READ 2 is done inside the DARAM4W, such as DARAM4W 113. The global traffic module 115 drives bank select signals only.

The M bus 121 write is always given top priority in the second half cycle. They will be serviced immediately and are never stalled. Program "writes" from the local subsystem are given next priority. Program "writes" will be stalled if an  
25 M bus 121 "write" request is asserted at the same time as a local program "write" request. READ 3 and READ 4 compete for the memory in the second half cycle. The DARAM4W are wired such that READ 3 and READ 4 are driven from the same half of the chip. That is, DARAM4W 113 in subsystem A 101, has READ  
30 3 connected to subsystem A 101 and READ 4 connected to

subsystem B 102. This is done to provide the most distant cores adequate set-up time.

The priority amongst READ 3 and READ 4 changes every time READ 3 and READ 4 both request access to the memory on the 5 same cycle, and there are no other requesters. Initially READ 3 will have priority over Read 4. A single register bit controls the priority. If a request from both READ 3 and READ 4 occurs simultaneously, priority is given to READ 3. The READ 3 request will complete, and the READ 4 request will be 10 stalled one clock. The priority register will toggle, so at the next occurrence of a simultaneous request by READ 3 and READ 4, READ 4 will be given top priority. The priority only changes when there is a collision between READ 3 and READ 4 and there are no other requesters for the second half cycle.

15 The arbitration logic for the READ 3 and READ 4 busses is contained within the DARAM4W, such as DARAM4W 113 of subsystem 101. This is done because the arbitration for second half access is slightly more involved than that of first half and the requesting cores are physically close to the target 20 memory. The multiplexing of the addresses and data for M bus 121 writes, program writes, READ 3 and READ 4 is done inside DARAM4W 113. Global traffic module 115 drives bank select signals 117 only.

The M bus 121 is driven by a local DMA controller and a 25 host port interface. Typically the M bus 121 will only request access to the SARAM 112 during initial program load. Under normal operating conditions, the M bus 121 will typically not access the DARAM4W 113. The program busses READ A, READ B, READ C, and READ D can be stalled for more than one wait state 30 if there is M bus 121 activity. If there is no M bus 121

activity, then the program READ busses will be stalled for one wait state at most.

Memory accesses through the peripheral port must be in the synchronous shared access mode (SAM). In shared access mode, the dual access RAM is accessible to both the DSP core and the peripheral. In this mode the peripheral accesses presented to the dual access RAM must be synchronous with the peripheral clock (slave). Asynchronous peripheral accesses are synchronized internally by the peripheral, and in case of a conflict between DSP and the peripheral, the peripheral has access priority and DSP access is delayed one clock cycle. The DSP accesses can only occur in SAM and are always synchronous with the DSP peripheral clock (slave).

A program read access could be stalled for one half of the cycle, while the second half of the cycle is not even used. For example, suppose only program reads 1 and 2 made requests to access the memory. Program access 1 could occur in the first half of the cycle, and 2 would be stalled one clock. No access will occur during the second half of the cycle. Note reduction of complexity in the arbitration results from permitting this kind of unused memory access slot.

To minimize the number of four-way shared memory data ports on the traffic module, the "read" data from the four-way shared memory banks is driven on to a single tri-state bus. The selects generated from the respective global traffic modules are used to control tri-state buffers.

Figure 3 illustrates conceptually the flow of data arbitrated within a subsystem. Subsystem A 101 is used as an example. Six request inputs are shown representing the six accesses which are arbitrated. Request 314 is associated with

an address "P1 Address" and request 315 is associated with an address "P3 Address". Four other similar requests can be simultaneously present at arbitration request inputs 330.

Arbitration and data steering logic 304 receives these inputs 5 and separate write data inputs from M bus 121 and E bus 122. Addresses 327 are sent to address steering logic 303. Address steering logic 303 supplies two addresses to multiplexer 326. Multiplexer 326 selected one address as controlled by strobe (STRB) signal 307. The selected address input A 317 contains 10 the required address for each half-clock cycle switched by multiplexer 326 as driven by STRB signal 307. STRB signal 307 and inverted opposite phase signal STRBZ (which are collectively labeled STRB 307) are derived in buffered form from the main DSP clock.

15 The DARAM 113 read port includes two full-word registers 301 and 302 which are clocked on opposite phases of SLAVE signal 311, which is a buffered form of the main DSP clock. Data Q 300 from the DARAM 113 is latched in the first phase of SLAVE signal 311 into register 301 and in the second phase of 20 SLAVE signal 311 into register 302. This allows P1 data 328 to arrive at the beginning of the first half of SLAVE signal 311 cycle and P3 data 329 to arrive at the beginning of the second half of SLAVE signal 311 cycle.

Blocks 305, 306, 320, and 325 provide bus switching.

25 Blocks 305, 306, 320 and 325 are controlled from arbitration and data steering logic 304 via SLAVE signal 311, control signal 312 and control signal 313, respectively. The example block diagram of Figure 3 could be modified in possible implementations. It is generally preferable to locate bus

switching outside of the individual subsystems as illustrated in Figure 1.

Figure 4 illustrates the RAM access timing for first-half arbitration that occurs between the two furthest subsystems. 5 Subsystem A 101 is once again used as an example. In Figure 4 the signal P1SEL 400 is generated as part of the arbitration algorithm, address 317 and data Q 300 are the address input and data output, respectively, from DARAM 113. Referring to Table 3, program read C and program read D would arbitrate for 10 subsystem A 101 DARAM4W 113 in the first half cycle arbitration. The P1 address from program read C and program read D is valid on P1 address bus 314 during the both phases 401 and 402 of the first clock cycle of SLAVE signal 311. The program bus is arbitrated and the winning address is presented 15 to the subsystem A 101 DARAM 113 on address bus 317 when the STRB signal is '0' at time 404. The P1 read data 328 from subsystem A 101 DARAM4W 113 is available during the next full clock cycle at phases 407, 408.

Figure 5 illustrates the RAM access timing 20 for second-half arbitration that occurs between the two closest subsystems. Subsystem A 101 is once again used as an example. Referring to Table 3, program A and program B arbitrate for memory in the second half arbitration. The address from program A and program B is valid on P3 address bus 315 during the first half-cycle 501, 502 of SLAVE signal 307. The program bus is arbitrated and the winning address is presented to the subsystem A 101 DARAM4W 113 on address bus 317 when the STRBZ signal 307 is "0". Note STRBZ signal 307 is "0" during the first half of SLAVE cycle 501, in contrast 25 to STRB of Figure 4 which was "0" during the second half of 30

TI-30704

4/9/01

the SLAVE cycle 402. The P3 read data 329 from DARAM4W 113 is available during the next SLAVE cycle 507, 508.

TELETYPE - DECODE